

**WHAT IS CLAIMED IS:**

1. A database system capable of executing a database application that transfers a logical object in multiple fragments, the database system comprising:
  - a main storage site;
  - a remote storage site capable of linking to the main storage site and mirroring information stored in the main storage site, the remote storage site including a storage and a cache sidefile divided into a plurality of array sidefile recordsets;
  - a main protocol executable on the main storage site and capable of including information indicative of logical object fragment commencement and completion in the multiple fragment database application transfer; and
  - a remote protocol executable on the remote storage site and capable of controlling the cache sidefile to cache the multiple fragments as received and to enable destaging of the logical object to the storage on receipt of all fragments.
2. The database system according to Claim 1 wherein:
  - the main protocol includes information indicative of logical object fragment commencement and completion using a technique selected from among a group including: (1) explicitly identifying starting and ending fragments, and (2) implicitly indicating either of the starting fragment and the ending fragment.
3. The database system according to Claim 1 wherein the main protocol further comprises:
  - an address translation process that translates a logical address to a list of physical addresses.
4. The database system according to Claim 1 wherein the main protocol further comprises:
  - an address translation process that resolves a virtual write address of the database application into a pick list of actual physical media writes associated with the logical object.

5. The database system according to Claim 1 wherein the main protocol further comprises:

a process capable of creating a control message for communication to the remote protocol that instructs individual physical storage elements to operate on the multiple physical writes as a single object entity so that all or none is destaged to the storage.

6. The database system according to Claim 1 wherein the main protocol further comprises:

a process capable of receiving an application request to write the logical object of a specified length to a specified virtualized storage address;  
a process capable of converting the virtualized write address and resolving the transfer length to designate at least one physical address in at least one physical storage device for transferring the logical object in fragments;  
a process capable of sending a first control message to the at least one physical storage device that delineates the start of a logical object that is to be held in a remote mirror cache for destaging; and  
a process capable of sending a second control message that delineates the end of the logical object so that the mirror cache is destaged to the at least one physical storage device, no portion of the logical object fragments being otherwise destaged.

7. The database system according to Claim 1 wherein:  
information is replicated from the main storage site to the remote storage site using a technique selected from among a group including: (1) synchronous data replication and (2) asynchronous data replication.

8. The database system according to Claim 1 wherein:  
the logical object multiple fragments are controllably destaged in all-or-none fashion to all devices in a consistency group.

9. An article of manufacture comprising:  
a controller usable medium having a computable readable program code embodied therein for executing in a database system that runs a database application for mirroring a logical object in multiple fragments from a main storage site to a remote storage site, the computable readable program code further comprising:  
a code capable of causing the controller to interface with the database application that links and mirrors data between the main storage site and the remote storage site, the remote storage site including a storage and a cache sidefile divided into a plurality of array sidefile recordsets; and  
a code capable of causing the controller to create and deploy control information indicative of logical object fragment commencement and completion in the multiple fragment database application transfer; the control information controlling the cache sidefile to cache the multiple fragments as received and to enable destaging of the logical object to the storage on receipt of all fragments.
10. The article of manufacture according to Claim 9 wherein the computable readable program code further comprises:  
a code capable of creating control information indicative of logical object fragment commencement and completion using a technique selected from among a group including: (1) explicitly identifying starting and ending fragments, and (2) implicitly indicating either of the starting fragment and the ending fragment.
11. The article of manufacture according to Claim 9 wherein the computable readable program code further comprises:  
a code capable of translating a logical address to a list of physical addresses.

12. The article of manufacture according to Claim 9 wherein the computable readable program code further comprises:

a code capable of resolving a virtual write address of the database application into a pick list of actual physical media writes associated with the logical object.

13. The article of manufacture according to Claim 9 wherein the computable readable program code further comprises:

a code capable of creating a control message for communication to the remote protocol that instructs individual physical storage elements to operate on the multiple physical writes as a single object entity so that all or none is destaged to the storage.

14. The article of manufacture according to Claim 9 wherein the computable readable program code further comprises:

a code capable of receiving an application request to write the logical object of a specified length to a specified virtualized storage address;  
a code capable of converting the virtualized write address and resolving the transfer length to designate at least one physical address in at least one physical storage device for transferring the logical object in fragments;  
a code capable of sending a first control message to the at least one physical storage device that delineates the start of a logical object that is to be held in a remote mirror cache for destaging; and  
a code capable of sending a second control message that delineates the end of the logical object so that the mirror cache is destaged to the at least one physical storage device, no portion of the logical object fragments being otherwise destaged.

15. The article of manufacture according to Claim 9 wherein the computable readable program code further comprises:

a code capable of replicating information from the main storage site to the remote storage site using a technique selected from among a group including: (1) synchronous data replication and (2) asynchronous data replication.

16. The article of manufacture according to Claim 9 wherein the computable readable program code further comprises:

a code capable of controllably destaging the logical object multiple fragments in all-or-none fashion to all devices in a consistency group.

17. An article of manufacture comprising:

a controller usable medium having a computable readable program code embodied therein for executing in a database system that runs a database application for mirroring a logical object in multiple fragments from a main storage site to a remote storage site, the computable readable program code further comprising:

a code capable of causing the controller to receive the logical object in multiple fragment transfers in combination with control information indicative of logical object fragment commencement and completion;

a code capable of causing the controller to control storage of the logical object multiple fragments in a cache sidefile divided into a plurality of array sidefile recordsets; and

a code capable of causing the controller to cache the multiple fragments as received and to enable destaging of the logical object to the storage on receipt of all fragments.

18. The article of manufacture according to Claim 17 wherein the computable readable program code further comprises:

a code capable of determining logical object fragment commencement and completion using a technique selected from among a group including: (1) receiving explicitly identified starting and ending fragments, and (2) deriving either of the starting fragment and the ending fragment implicitly from received control information.

19. The article of manufacture according to Claim 17 wherein the computable readable program code further comprises:

a code capable of causing the controller to track order of fragment updating between the main storage site and the remote storage site including updating of the sidefile recordsets.

20. The article of manufacture according to Claim 17 wherein the computable readable program code further comprises:

a code capable of replicating information from the main storage site to the remote storage site using a technique selected from among a group including: (1) synchronous data replication and (2) asynchronous data replication.

21. The article of manufacture according to Claim 17 wherein the computable readable program code further comprises:

a code capable of controllably destaging the logical object multiple fragments in all-or-none fashion to all devices in a consistency group.

22. A storage element readable by a controller tangibly embodying a program of instructions executable by the controller to perform method acts for executing in a database system that runs a database application for mirroring a logical object in multiple fragments from a main storage site to a remote storage site, the method acts comprising:

receiving the logical object in multiple fragment transfers in combination with control information indicative of logical object fragment commencement and completion;

controlling storage of the logical object multiple fragments in a cache sidefile divided into a plurality of array sidefile recordsets; and

caching the multiple fragments as received and to enable destaging of the logical object to the storage on receipt of all fragments.

23. The storage element according to Claim 22 wherein the method acts further comprise:

creating control information indicative of logical object fragment commencement and completion using a technique selected from among a group including: (1) explicitly identifying starting and ending fragments, and (2) implicitly indicating either of the starting fragment and the ending fragment.

24. The storage element according to Claim 22 wherein the method acts further comprise:

resolving a virtual write address of the database application into a pick list of actual physical media writes associated with the logical object.

25. The storage element according to Claim 22 wherein the method acts further comprise:

creating a control message for communication to the remote protocol that instructs individual physical storage elements to operate on the multiple physical writes as a single object entity so that all or none is destaged to the storage.

26. The storage element according to Claim 22 wherein the method acts further comprise:

receiving an application request to write the logical object of a specified length to a specified virtualized storage address;

converting the virtualized write address and resolving the transfer length to designate at least one physical address in at least one physical storage device for transferring the logical object in fragments;

sending a first control message to the at least one physical storage device that delineates the start of a logical object that is to be held in a remote mirror cache for destaging; and

sending a second control message that delineates the end of the logical object so that the mirror cache is destaged to the at least one physical storage device, no portion of the logical object fragments being otherwise destaged.

27. The storage element according to Claim 22 wherein the method acts further comprise:

replicating information from the main storage site to the remote storage site using a technique selected from among a group including: (1) synchronous data replication and (2) asynchronous data replication.

28. A storage element readable by a controller tangibly embodying a program of instructions executable by the controller to perform method acts for executing in a database system that runs a database application for mirroring a logical object in multiple fragments from a main storage site to a remote storage site, the method acts comprising:

interfacing with the database application that links and mirrors data between the main storage site and the remote storage site, the remote storage site including a storage and a cache sidefile divided into a plurality of array sidefile recordsets; and

creating and deploying control information indicative of logical object fragment commencement and completion in the multiple fragment database application transfer; the control information controlling the cache sidefile to cache the multiple fragments as received and to enable destaging of the logical object to the storage on receipt of all fragments.

29. The storage element according to Claim 28 wherein the method acts further comprise:

determining logical object fragment commencement and completion using a technique selected from among a group including: (1) receiving explicitly identified starting and ending fragments, and (2) deriving either of the starting fragment and the ending fragment implicitly from received control information.

30. The storage element according to Claim 28 wherein the method acts further comprise:

tracking order of fragment updating between the main storage site and the remote storage site including updating of the sidefile recordsets.



31. The storage element according to Claim 28 wherein the method acts further comprise:  
replicating information from the main storage site to the remote storage site using a technique selected from among a group including: (1) synchronous data replication and (2) asynchronous data replication.
32. The storage element according to Claim 28 wherein the method acts further comprise:  
controllably destaging the logical object multiple fragments in all-or-none fashion to all devices in a consistency group.